

09-29-00

A

FISH & RICHARDSON P.C.,P.A.

3300 Dain Rauscher Plaza
60 South Sixth Street
Minneapolis, Minnesota
55402

Telephone
612 335-5070

Facsimile
612 288-9696

Web Site
www.fr.com

September 28, 2000

Attorney Docket No.: 10559/364001/P8247-2

Box Patent Application
Commissioner for Patents
Washington, DC 20231

Presented for filing is a new continuation-in-part patent application of:

Applicant: DAVID I. POISNER

Title: ACCESSING MULTI-PORTED MEMORY

Enclosed are the following papers, including those required to receive a filing date
under 37 CFR §1.53(b):

	<u>Pages</u>
Specification	9
Claims	6
Abstract	1
Declaration	[To be Filed at a Later Date]
Drawing(s)	3

Enclosures:
— Postcard.

There are 30 total claims, 6 of which are independent.

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL419343051US

I hereby certify that this correspondence is being deposited with the
United States Postal Service as Express Mail Post Office to Addressee
with sufficient postage on the date indicated below and is addressed to
the Commissioner for Patents, Washington, D.C. 20231.

September 28, 2000
Date of Deposit

Signature

Vince Defante
Typed or Printed Name of Person Signing Certificate

09/28/00
Jc929 U.S. PTO
Frederick P. Fish
1855-1930
W.K. Richardson
1859-1951

Jc836 U.S. PTO
09/672345
09/28/00

FR
BOSTON
DALLAS
DELAWARE
NEW YORK
SAN DIEGO
SILICON VALLEY
TWIN CITIES
WASHINGTON, DC

Commissioner for Patents
September 28, 2000
Page 2

This application is a continuation-in-part (and claims the benefit of priority under 35 USC 120) of U.S. application serial no. 09/572,047, filed May 16, 2000. The disclosure of the prior application is considered part of (and is incorporated by reference in) the disclosure of this application.

Basic filing fee	\$0
Total claims in excess of 20 times \$18	\$0
Independent claims in excess of 3 times \$78	\$0
Fee for multiple dependent claims	\$0
Total filing fee:	\$0

Under 37 CFR §1.53(f), no filing fee is being paid at this time.

If this application is found to be incomplete, or if a telephone conference would otherwise be helpful, please call the undersigned at (858) 678-5070.

Kindly acknowledge receipt of this application by returning the enclosed postcard.

Please send all correspondence to:

SCOTT C. HARRIS
Fish & Richardson P.C.
Customer Number: 20985
4350 La Jolla Village Drive, Suite 500
San Diego, California 92122

Respectfully submitted,



Daniel J. Hanson
Reg. No. 46,757
Enclosures
DZH/jaw
60022850.doc

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: ACCESSING MULTI-PORTED MEMORY

APPLICANT: DAVID I. POISNER

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL419343051US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

Date of Deposit September 28, 2000

Signature 

Vince Defante
Typed or Printed Name of Person Signing Certificate

ACCESSING MULTI-PORTED MEMORY

This is a continuation in part of Ser. No. 09/572,047,
filed May 16, 2000.

BACKGROUND

5 This invention relates to accessing multi-ported memory.

In a conventional computing system the central processing
unit (CPU), main memory and input/output (I/O) devices are
connected by a bus. A "bus master" or "bus arbiter" controls
and directs data traffic among the components of the computing
10 system.

Main memory is used as the principal site for storing
data. An "access" to main memory writes data to or reads data
from main memory. Making an access (or "accessing") is
typically preceded by a request for access from another
15 component of the system, such as the CPU or an I/O device,
followed by a grant of permission by the bus arbiter.

There are two principal types of accesses. The first
type is a data access, in which large amounts of data are
written to or read from main memory. A data access may be on
20 the order of thousands of bytes. The second type is a
control/status access, characterized by a small number of
reads or writes to a defined data structure in order to report
the status of an input/output device, process data, or

initiate some input/output activity. In contrast to data accesses, a control/status access is usually on the order of a few bits. Control accesses are generally initiated by the CPU, while status accesses are generally initiated by the I/O devices.

DESCRIPTION OF DRAWINGS

Figs. 1 and 2 are conceptual block diagrams depicting an embodiment of the invention.

Fig. 3 is a conceptual block diagram depicting a component shown in Fig. 2, illustrating an embodiment of the invention.

DETAILED DESCRIPTION

In Fig. 1, a computer architecture 10 includes a central processing unit (CPU) 12, main memory 22 and I/O devices 26.

Main memory 22 is generally a form of random access memory (RAM), such as dynamic RAM (DRAM), Rambus™ DRAM (RDRAM), synchronous DRAM (SDRAM) or SyncLink DRAM (SLDRAM).

Communications among these components are regulated by a memory controller 16. Memory controller 16 performs the functions of a bus arbiter by managing communications 14 with CPU 12, communications 20 with main memory 22, and communications 24 with I/O devices 26.

Memory controller 16 includes a dual-ported memory 18 that can be accessed by CPU 12 and by I/O device 26 independently. Dual-ported memory 18 can be any form of random access memory although static random access memory, or SRAM, is useful because of its speed and reliability, and some forms of DRAM are useful because DRAM saves chip space. In the embodiments described below, the memory is dual-ported, but the invention is not necessarily limited to two ports. Dual-ported memory 18 may be integrated into a single integrated circuit with memory controller 16, which improves speed and saves chip space.

Typically control/status accesses waste system resources because they are generally short and have comparatively few bits. Many forms of main memory, such as RDRAM, handle long strings of bits more efficiently than short strings of bits. In addition, as will be discussed below, many conventional computer architectures include a cache memory and ensure cache consistency by employing "bus snooping," i.e., monitoring the communication channels for write accesses to main memory. Bus snooping may require additional logic to be present in each microprocessor and, therefore, can increase the cost and decrease the performance of the microprocessor. The performance degradation incurred by bus snooping is

accelerated as the number of short, control/status accesses increases relative to the number of data accesses.

Computer architecture 10 offers a more efficient use of resources by directing control/status accesses to dual-ported memory 18 in controller 16, rather than to main memory 22. By directing control/status accesses to dual-ported memory 18 instead of main memory 22, a bus snoop and its attendant delay are avoided.

There are several ways in which accesses from CPU 12 and I/O device 26 may be directed to dual-ported memory 18. One way is for the computer's operating system to direct appropriate accesses to dual-ported memory 18. Another way is for components to be configured to look for dual-ported memory 18. For example, a controller for I/O device 26 may check controller 16 to see whether dual-ported memory 18 is available, and if dual-ported memory 18 is available, I/O device 26 may use dual-ported memory 18. If dual-ported memory 18 is not available, I/O device 26 may use main memory 22. Other ways in which accesses may be directed to controller's memory exist, and these are not exclusive of each other.

Fig. 2 is a block diagram depicting a computer architecture 40 that includes an embodiment of the invention. Fig. 2 is more detailed than Fig. 1, and shows more of the

components typically present within a computer architecture than are shown in Fig. 1. Architecture 40 includes main memory 46 electrically coupled to memory controller 50. Architecture 40 further comprises CPU 42, which is

5 electrically coupled to memory controller 50 and to cache memory 44. Cache memory 44 comprises high-speed RAM for holding data and instructions that, for example, are contained in main memory 46, or originate from main memory 46, and that are recently used by CPU 42. For cache memory 44 to be

10 useful, the information stored in cache memory 44 must be consistent with the data and instructions held in main memory 46. As previously mentioned, bus snooping is a technique used to maintain cache consistency. Cache memory 44 can copy data and instructions from main memory 46 via memory controller 50.

15 In addition, memory controller 50 is coupled to I/O bus 54. Other components may be electrically coupled to memory controller 50 in addition to those shown in Fig. 2.

I/O bus 54 is a communication channel typically dedicated to input/output operations. I/O bus 54 may be a high-speed

20 bus supporting peripherals operating at high data throughput rates, such as a Peripheral Component Interconnect (PCI) bus. The term "peripheral" encompasses but is not limited to external devices attached to the computer system. I/O devices are the most common peripherals. A computer system may have

many peripherals, and these devices may interface with the system by way of I/O bus 54.

In Fig. 2, three peripherals are shown for exemplary purposes: an external storage device 62 such as a hard disk, a display device 64 such as a monitor, and a printer 66 such as a laser printer. The peripherals 62, 64, 66 interface with the system through I/O controllers 56, 58, 60, respectively. Although three peripherals 62, 64, 66 are shown, I/O bus 54 is not necessarily limited to use by three peripherals. Other peripherals include devices such as a modem, keyboard, optical drive, network connection, and mouse.

Memory controller 50 also provides a communication connection among components such as CPU 42 and main memory 46. For example, memory controller 50 directs data traffic among CPU 42, main memory 46 and I/O bus 54. Bus arbiter 48, included in memory controller 50, regulates data traffic on I/O bus 54. In addition, memory controller 50 includes dual-ported memory 52, independently accessible by CPU 42 or by any peripheral 62, 64, 66 via I/O bus 54. More specifically, the operating system executing on CPU 42 treats dual-port memory 52 as non-cached memory. As a result, read and write accesses to dual-ported memory 52 do not implicate cache memory 44 and need not be monitored. Therefore, dual-ported memory 52 may

be accessed by CPU 42 or by any peripheral 62, 64, 66 without a bus snoop, avoiding the delay caused by the bus snoop.

In this architecture 40, control/status accesses are directed principally to dual-ported memory 52, rather than to main memory 46, thereby avoiding a bus snoop and its attendant delay. The advantages described above, such as speed and efficiency, can be attained with this architecture 40, and "smart" addressing can be used to map a virtual address to a physical address. With smart addressing, addresses are mapped to simplify conversion of the virtual address to the physical address. Smart addressing can avoid operating system calls and can save time by relating virtual and physical addresses through a simple mapping algorithm. Byte-aligning addresses of a page boundary on an external disk 62 with addresses in dual-ported memory 52 is one example of smart addressing. With the addresses byte-aligned, the addresses are related by the addition or subtraction of a constant. With such a mapping, it is not necessary for a device driver to call the operating system to convert a virtual address to a physical address; the device driver can translate the virtual address to a physical address simply by adding or subtracting.

As is evident from Fig. 2, increased usage of dual-ported memory 52 by peripherals results in reduced usage of main memory 46 and consequently less data traffic on the channel 45

connecting main memory 46 to memory controller 50. With less traffic on channel 45, CPU 42 can have greater access to main memory 46, and more information can be transferred between CPU 42 and main memory 46 in a given period of time.

5 Memory controller's dual-ported memory 52 need not be restricted to reading or writing control/status accesses. In some circumstances, data may be stored in memory 52.

As shown in Fig. 3, dual-ported memory 52 can be divided into general-purpose memory 80 and reservation bits memory 82. General-purpose memory 80 may be subdivided into regions or
10 blocks of memory, which may be allocated for use by CPU 42 or by I/O device 62, 64 and 66. Each block may be uniquely identified with a reservation bit, such that for N blocks there are N corresponding reservation bits. For example, the
15 first reservation bit 84 is identified with the first block 86. The number of required reservation bits depends upon the size of general-purpose memory 80 and upon the size of a block. For example, if a block consists of 128 bytes, and if general-purpose memory 80 includes 4 kilobytes of memory, then
20 32 blocks would be available and 32 reservation bits would be needed (32 bits x 128 bytes/bit = 4 kilobytes).

The state of a reservation bit denotes whether the reservation bit's corresponding memory block in general-purpose memory 80 has been allocated. For example, a '0'

value bit in the first reservation bit 84 may signify that the first memory block 86 is free, while a '1' value bit in the first reservation bit 84 may signify that the first memory block 86 has been allocated for use. When a block of general-purpose memory 80 in dual-ported memory 52 is needed, the reservation bits 82 corresponding to the blocks can be checked to determine whether any blocks are free. When a block in general-purpose memory 80 has been allocated but is no longer needed, the block can be de-allocated by clearing the corresponding reservation bit to indicate that the block is free. In the event all reservation bits 82 show that all general-purpose memory 80 has been allocated, then main memory 48 may be accessed instead.

Fig. 3 exemplifies one of many possible memory organizations in dual-ported memory 52. Reservation bits 82 need not be grouped together as shown in Fig. 3, but may, for example, be denoted as the initial bit of each block. Furthermore, each block of general-purpose memory may have multiple reservation bits mapped to the block.

Other embodiments are within the scope of the following claims.

WHAT IS CLAIMED IS:

- 1 1. A computer system, comprising:
2 a non-cached multi-ported memory;
3 a central processing unit coupled to the multi-ported
4 memory;
5 a peripheral device coupled to the multi-ported memory;
6 the central processing unit and the peripheral device
7 being configured to access the multi-ported memory
8 independently.
- 1 2. The system of claim 1, further comprising an
2 operating system executing on the central processing unit,
3 wherein the operating system is configured such that accesses
4 to the multi-ported memory are not cached.
- 1 3. The system of claim 1, wherein the multi-ported
2 memory is dual-ported.
- 1 4. The system of claim 1, wherein the multi-ported
2 memory is embedded within a memory controller.
- 1 5. The system of claim 4, wherein the multi-ported
2 memory and memory controller are integrated into a single
3 chip.

1 6. The system of claim 1, wherein the multi-ported
2 memory is static random access memory or dynamic random access
3 memory.

1 7. The system of claim 1, wherein the multi-ported
2 memory stores reservation bits mapped to blocks of general-
3 purpose memory in the multi-ported memory.

1 8. The system of claim 1, wherein virtual addresses
2 within multi-ported memory are mapped to physical addresses
3 with smart addressing.

1 9. The system of claim 1, wherein the coupling of the
2 peripheral device to the memory controller includes an
3 input/output bus.

1 10. A method comprising:
2 making, from a peripheral device, a data access to memory
3 in a computer;
4 making, from the peripheral device, a status access to
5 memory in the computer;
6 routing the data access to a first memory in the
7 computer; and

8 routing the status access to a second memory in the
9 computer.

1 11. The method of claim 10, wherein the first memory
2 comprises main memory.

1 12. The method of claim 10, wherein the second memory
2 comprises memory included in a memory controller.

1 13. The method of claim 10, wherein the second memory is
2 dual-ported.

1 14. An article comprising a computer-readable medium
2 which stores computer-executable instructions for memory
3 accessing, the instructions causing a machine to:

4 make, from a peripheral device, a data access to memory
5 in a computer;

6 make, from the peripheral device, a status access to
7 memory in the computer;

8 route the data access to a first memory in the computer;
9 and

10 route the status access to a second memory in the
11 computer.

1 15. The article of claim 14, wherein the computer
2 includes an input/output controller.

1 16. The article of claim 14, wherein the first memory
2 comprises main memory.

1 17. The article of claim 14, wherein the second memory
2 comprises memory included in a memory controller.

1 18. The product of claim 14, wherein the second memory
2 is dual-ported.

1 19. A method comprising:
2 making, from a central processing unit, a data access to
3 memory in a computer;
4 making, from a peripheral device, a control access to
5 memory in the computer;
6 routing the data access to a first memory in the
7 computer; and
8 routing the control access to a second memory in the
9 computer.

1 20. The method of claim 19, wherein the first memory
2 comprises main memory.

1 21. The method of claim 19, wherein the second memory
2 comprises memory included in a memory controller.

1 22. The method of claim 19, wherein the second memory is
2 dual-ported.

1 23. An article comprising a computer-readable medium
2 which stores computer-executable instructions for memory
3 accessing, the instructions causing a machine to:

4 make, from a central processing unit, a data access to
5 memory in a computer;

6 make, from the central processing unit, a control access
7 to memory in the computer;

8 route the data access to a first memory in the computer;
9 and

10 route the control access to a second memory in the
11 computer.

1 24. The article of claim 23, wherein the first memory
2 comprises main memory.

1 25. The article of claim 23, wherein the second memory
2 comprises memory included in a memory controller.

1 26. The article of claim 23, wherein the second memory
2 is dual-ported.

1 27. An integrated circuit comprising:
2 a memory controller including at least two electrical
3 ports for coupling to communication channels; and
4 multi-ported memory communicatively coupled to each port.

1 28. The integrated circuit of claim 27, wherein the
2 multi-ported memory is dual-ported.

1 29. The integrated circuit of claim 27, wherein the
2 multi-ported memory is static random access memory or dynamic
3 random access memory.

1 30. The integrated circuit of claim 27, wherein the
2 multi-ported memory stores reservation bits mapped to blocks
3 of general-purpose memory in the multi-ported memory.

ACCESSING MULTI-PORTED MEMORY

ABSTRACT OF THE DISCLOSURE

A computer system comprising multi-ported memory electrically coupled to a central processing unit and a peripheral device. The central processing unit and the peripheral device access the multi-ported memory independently.




Fig. 1

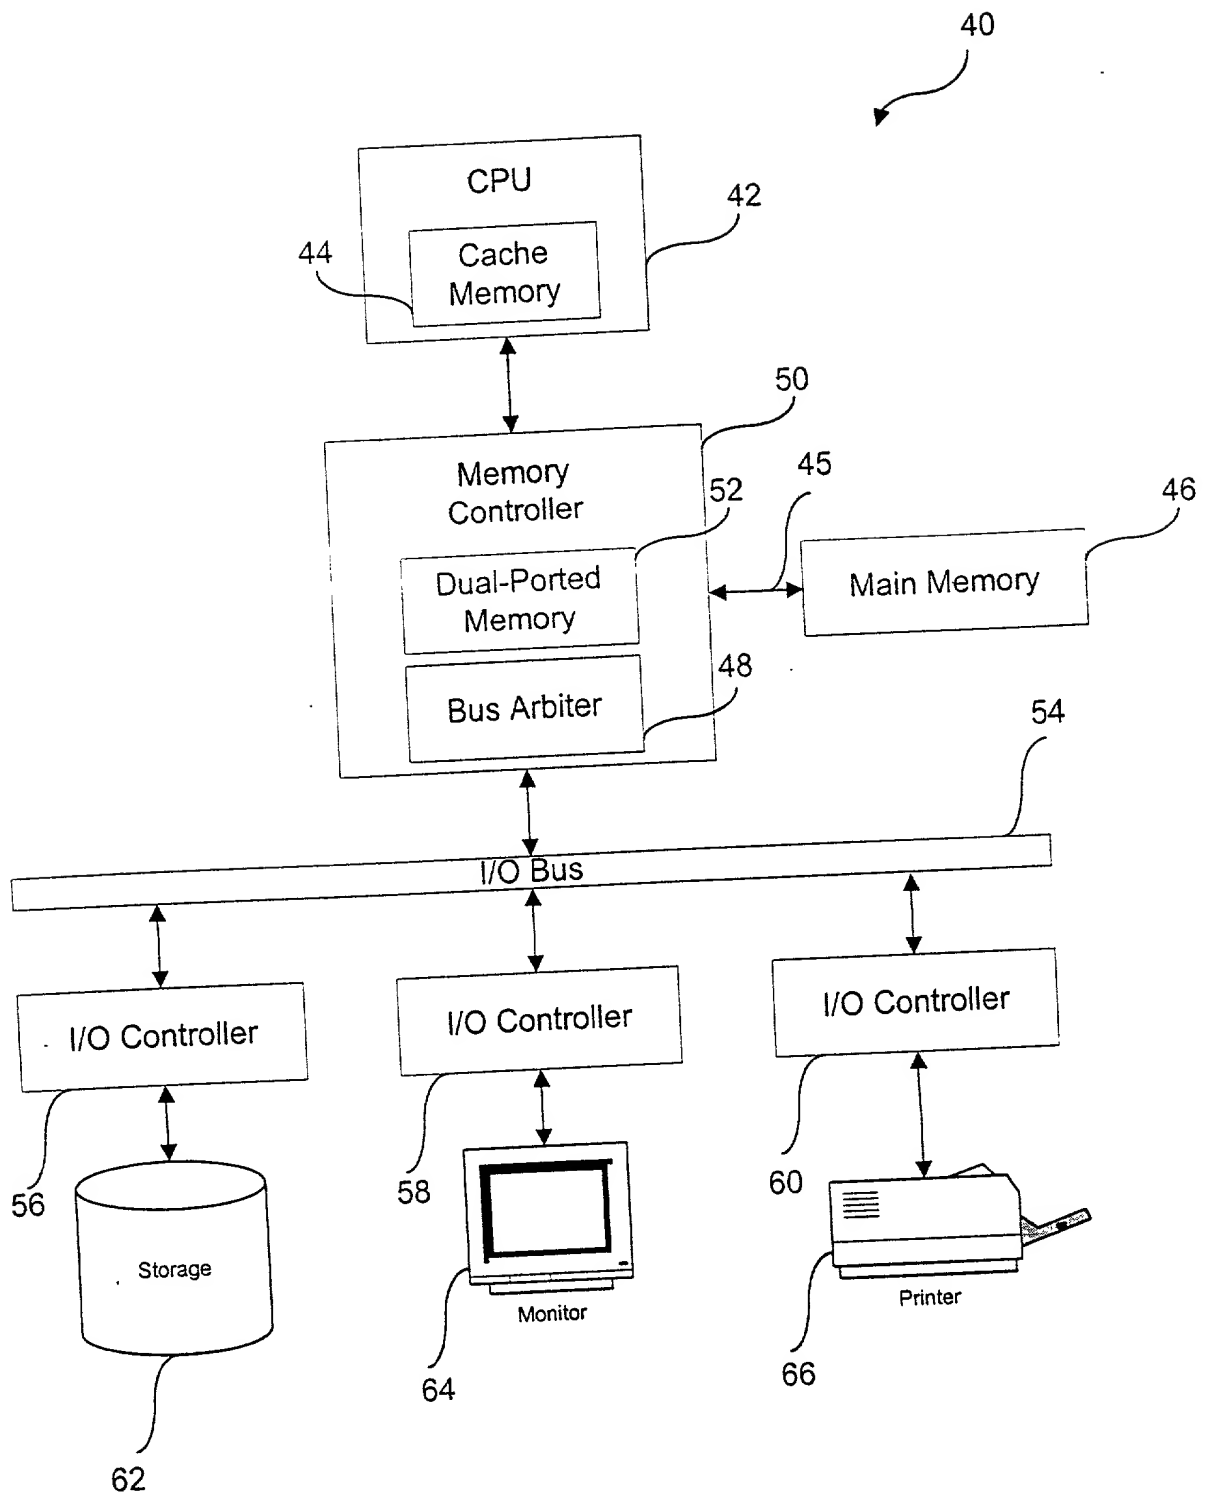


Fig. 2




Fig. 3